Lessons from the Learning Curve

Part 2 of Out of the Box: Overcoming Our Mind Set Legacies

Robert W. Hall

The phrase "learning curve" has come to mean almost any situation in which someone accumulates more knowledge of a business or process. However, few people except specialized cost analysts have formally used the learning curve model.

Developed by Dr. T.P. Wright in 1936, learning curves projected World War II production capacities. Afterward they became standard in the aircraft industry. In defense contracting, an expected learning rate is often factored into negotiations. However, learning curves seemingly made little impact on manufacturing progress in aviation or elsewhere. The history of the learning curve illustrates our 20th century mindset in North American manufacturing, a mindset whose time should have gone, but that is too deeply embedded to easily dislodge.

Manufacturing Progress

The learning curve was one of the earliest operations research models to emerge in World War II era, and widely used to predict war production output. Shortly after this tool's "buzz" peaked in the 1950s, a pair of experienced manufacturing researchers, R.W. Conway and Andrew Shultz, critiqued the use and validity of the learning curve in a well-known article, calling it the "manufacturing progress function."[1]

They saw a number of shortcomings. Because of the "mental box" they were in, they didn't see others that are more obvious today. All models have limitations, and some of their assumptions are easily seen. Others are too much a part of us to stand out.

Conway and Shultz noted that very few performance measures were tracked. A two-dimensional graph can plot only one variable against output, it does not suggest that manufacturing progress is multi-dimensional. Learning curves relate improvement to only one variable, cumulative production. In practice, most companies used the learning curve to predict either unit cost or labor hours per unit. (A brief review of the model is in the box copy.)

But the biggest problem was lack of data, or lack of timely data. In pre-computer days, summary data from the shop, particularly unit cost data, arrived as historical reports. In practice, most learning curve projections were extensions of cost systems.

In the 1950s cost systems were as cumbersome as today. Assumptions to assign costs and spread overhead were necessary, but arbitrary. Few cost systems track data to specific units produced, so these were imputed from averages, and analysts often had to "tease out" cost data. Even deciding how many units were complete was not clear-cut. Then as now, the devil plagued the details.

Production was frequently interrupted. Much of it was batch, not continuous. If the time between batches was so long that the same equipment and personnel were not used for subsequent runs, should that interruption be considered normal, or as a new start?

Some analysts plotted the learning against time, assuming that production was nearly linear with time. Conway and Shultz suggested that unit-specific data directly from the process source gave better results. Cost averages were too remote, hiding too many details, and creating artificially smooth curves.

Despite its many problems, Conway and Shultz's data showed that manufacturing progress according to learning curves appeared real. It wasn't mostly the spreading of fixed cost over more units as production accumulated, as alleged by a few critics. They assured managers that during ramp up, manufacturing inexorably continued to improve as more units were produced.

Conway and Shultz had no suggestions on how to make learning go faster. Figure 1 is a summary of learning rate factors as they saw them.

"Toe Up" and "Spurts"

At the time, the Conway and Shultz article was noted for revealing "toe up." Using data from several companies, they showed that some point learning appeared either to stop, or to slow considerably. Reality was that manufacturing improvement did not always progress without end.

Another issue was that analysts usually showed only nice, smooth graphs using cost data or other unit estimates from averages. However, data detailed unit by unit showed that improvement did not follow a smooth curve. Improvement came in spurts, sometimes interrupted by backsliding. Although they were not emphatic about it, Conway and Shultz saw that cost data, and managing by it, stays remote from reality.

Both of these phenomena are depicted in Figure 2. The erratic data did not surprise Conway and Shultz. They attributed "toe up" to nearing the end of ramp up or of a production contract, so that resources were diverted to other processes. Also, at some point the production process no longer painted targets of opportunity obvious to managerial radar.

Pre-Production Factors	**During-Production Factors**
Tooling: Type and degree of completion before start of production	Tooling changes during production for increasing capacity
Equipment and tool selection	Methods changes and simplification
Product design for manufacturability	Design changes: effectiveness of ECOs
Methods: effectiveness of pre-production preparation	Improvements in management and planning methods
State-of-the-art: Ability of the organization to cope with the level of technology	Volume changes: changes in rate or in anticipated duration of production that affect other decisions (like investment)
Magnitude of process design effort: attention to specs, test, inspections, etc.	Quality improvement: reduction of rework and repair
Shop organization: material handling, training, skills, planning, organization, etc.	Incentive pay plans: manner in which administered and when installed
	Operator learning or skill development (considered to be a very minor factor).

Summarized from "The Manufacturing Progress Function," R.W. Conway and Andrew Schultz, Jr., The Journal of Industrial Engineering, *Vol. 10, No. 1, Jan.-Feb. 1959, pp. 39-54.*

Figure 1. *Conway and Shultz's Learning Curve Improvement Factors (1959).*

Toe Up and Erratic Improvement Compared with a Smooth Learning Curve

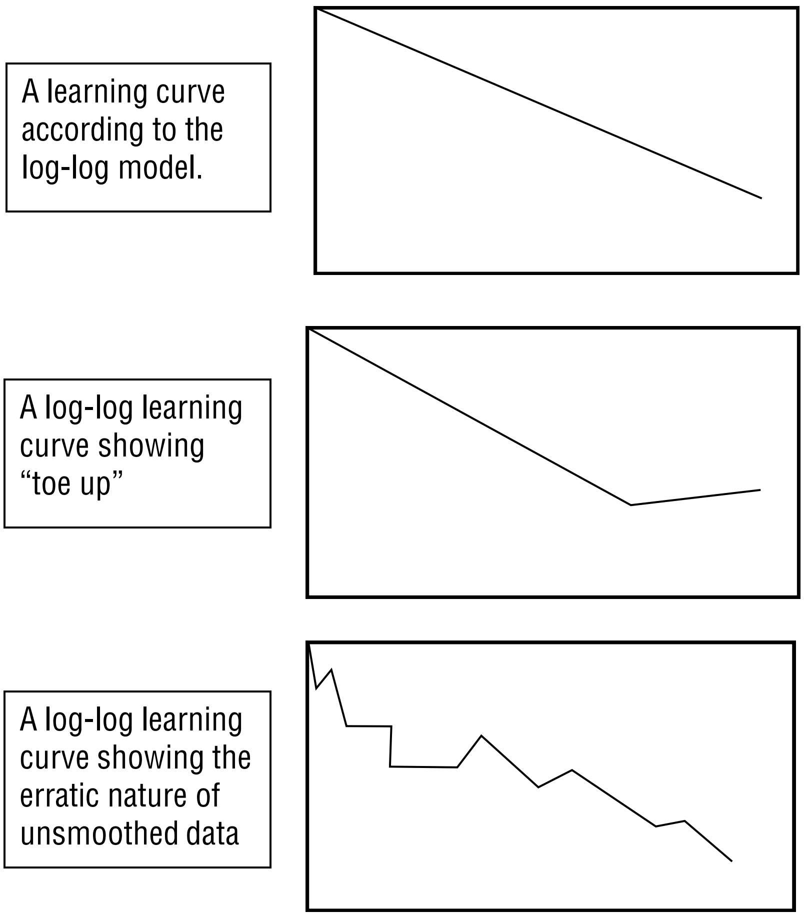

Figure 2.

Keeping Learning in a Box

Looking back on it, toe up indicated that process improvement was in a mental box constructed of managerial thinking.[2] At the time, few people had any cohesive concept of overall process improvement. It was the domain of engineers and managers, so improvement was largely "engineered," as revealed in Figure 1. Unless it was glaringly bad, improving workflow received little attention.

Leadtime, throughput time, or cycle time were not popular measures, but not totally unknown. As evidence, Figure 3 is a replica of a detailed learning curve of cycle time to assemble center wings of B-24 bombers at Willow Run. It shows the cycle time for each job from early 1942 to May 1943. (They also kept a learning curve on man-hours per unit.) The cycle time curve was discontinued on Job 600, when it hit 44 hours. After about 7800 more bombers were built, cycle time was down to 9 hours, so "learning" continued long after data collection for the curve ceased. However, by Job 600, it became obvious that capacity would be sufficient to meet the goal of a bomber an hour.

The literature of the time shows that the managerial mind set using learning curves limited its use mostly to ramp up. When production hit a budget goal or a capacity goal, improvement effort cut back substantially. At that point, the process was ready for "normal operation," or "turning over to production."

The behavioral side of this mind set is perfectly captured in the first sentence of another old article describing cost reduction methods: "Enactment of cost reduction programs is a management prerogative."[3]

No blather about worker participation. No drivel about teamwork. Management takes responsibility, getting all the credit, and sometimes the blame.

Anyone seriously pursuing lean manufacturing, much less other excellence goals, knows that this mind set is a serious blockage to progress, and one difficult to escape. In effect, we consider process improvement to be a program. When the chosen improvement initiatives are in place and the teams are functioning, we consider lean to be "operational." Usually many other issues then consume leadership attention. Mind sets have expanded one stage beyond ramp up, but learning is still closed in by a mental box as shown in Figure 4.

The Traditional Learning Curve Model

A learning curve is a graph showing that performance using some criterion changes by X percent each time production volume doubles. Since labor hours or unit cost have been by far the most popular performance criteria used, a learning curve is typically described as a percentage. For example, if unit cost decreases by 20 percent each time cumulative production doubles, it's called an 80 percent progress function. The steeper the learning curve, the faster the learning.

In this case, "learning" refers to process improvement – results, not mere insight. Since improvement is marked each time production doubles, this is a non-linear curve. A typical shape is shown below on a linear scale graph.

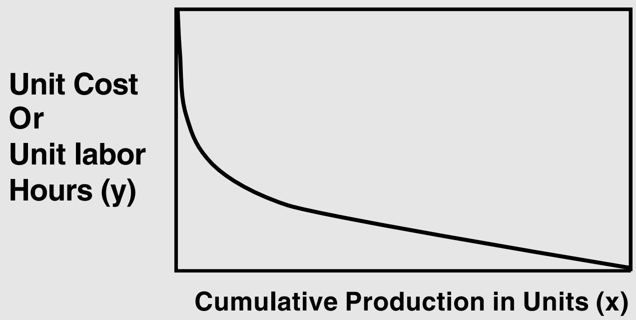

Improvement appears very steep at first. For example, consider reduction in unit cost by a 70 percent curve. The second unit's cost is 70 percent of the first; the fourth, 70 percent of the second, and so on. By the time production cumulates to around 100,000 units, one must make another 100,000 units to see another 70 percent reduction. The logic of this model makes "learning" appear to tail off.

The learning curve formula is:

$y_x = y_1 (x^{-b})$, where y_1 is the unit cost (or labor hours) for the first unit built.

However, a learning curve is usually shown in logarithmic form:

$(\log y_x) = (\log y_1) - b(\log x)$

On a log-log scale the "curves" appear as straight lines, a phenomenon more familiar to those old enough to have used slide rules, whose scales were based on logarithms.

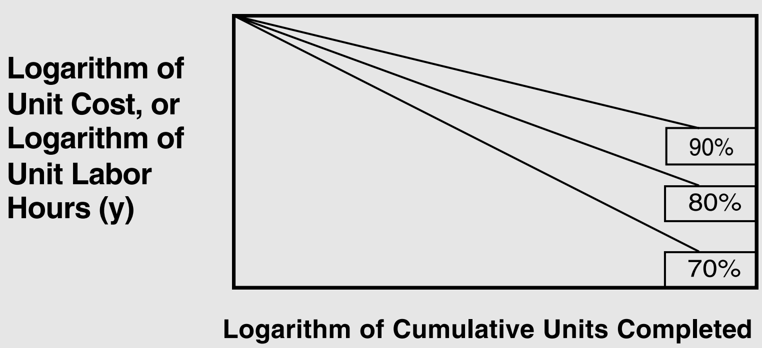

Besides tracking the current unit cost, the curve has been used to track cumulative cost, that is, how much has been spent for all production up through the "xth" unit. Dividing by the number of units thus far produced yields the cumulative average cost per unit. If you have taken on a fixed cost contract, that's a matter of considerable interest.

To the profit-minded, the difference in progress between a 70 percent curve and a 90 percent curve is the difference between making a bundle and losing your shirt. Over the past half-century, the learning curve has been used when bidding on contract production, especially in the defense business. If a bid assumed a 70 percent learning rate, but the process only improved at a 90 percent rate, disaster was likely. Even a two percent gap between an estimate and the actual is troublesome.

Cycle Time Learning Curve for B-24 Center Wings

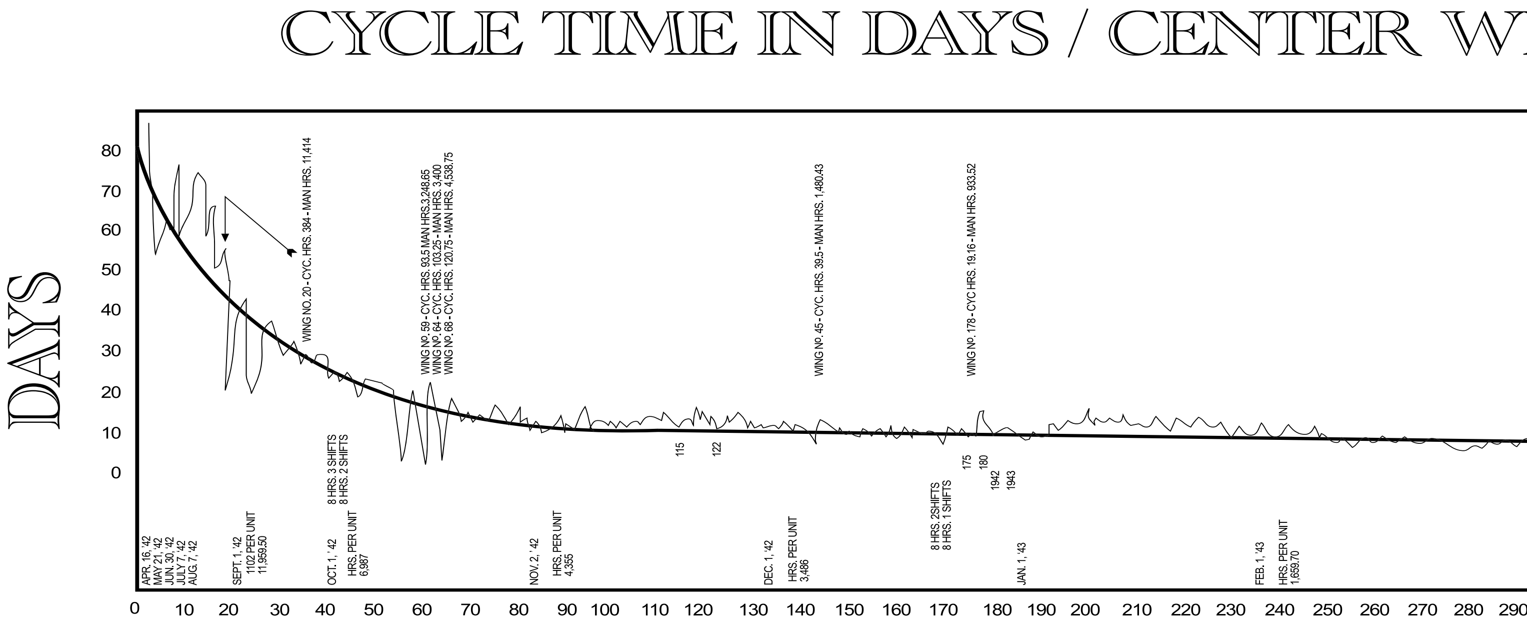

The above graph indicated how the cycle time required to assemble the Center Wing section was decreased during the building of the first 600, commencing in early 1942 and extending into May, 1943. The contrast may be illustrated by comparing the cycle hours for Wing No. 20 which was 384 hours with Wing 383 which was 44.33 hours. Another comparison: the man hours for Wing No. 20 were 11,414 while the March average was 1260.

Figure 3. *This curve is graphed on a linear scale, and is typical of other detailed curves. From the collections of the Research Center, Henry Ford Museum and Greenfield Village, Accession. 435, Box 40, Willow Run.*

The mental box limiting learning

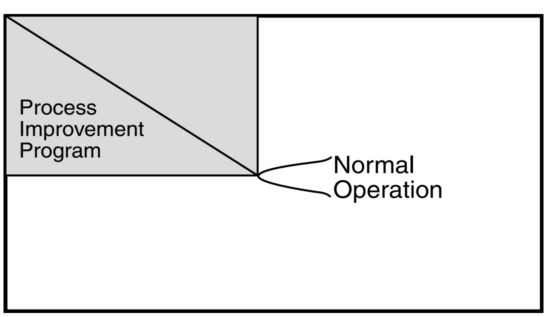

Figure 4.

Note the normal curve outside the learning curve box. It denotes that a process has now reached a goal, so that variations thereafter are more like deviations from an ongoing, stable process. By the logic of the normal curve, the objective becomes control, holding the process within its design limits, with minimal attention to further improvement.

Other Applications of the Power Law

The learning curve is but one example of a power law relationship. Almost any function that plots as a straight line on log-log paper follows a power law. Many basic phenomena have been plotted as a power law, for example the frequency of word usage in a language. That relationship, called Zipf's Law, is shown in Figure 5.

Word count frequency is used to evaluate whether two different pieces of writing are from the same author. If the word frequency of an unknown writing is similar enough to that of an authentic work, the chances are high that the same person wrote both. Many authors have "signature" word frequencies just as by ear a characteristic style of music is considered likely to have

ING VERTICAL FIXTURES / DEPT. 936

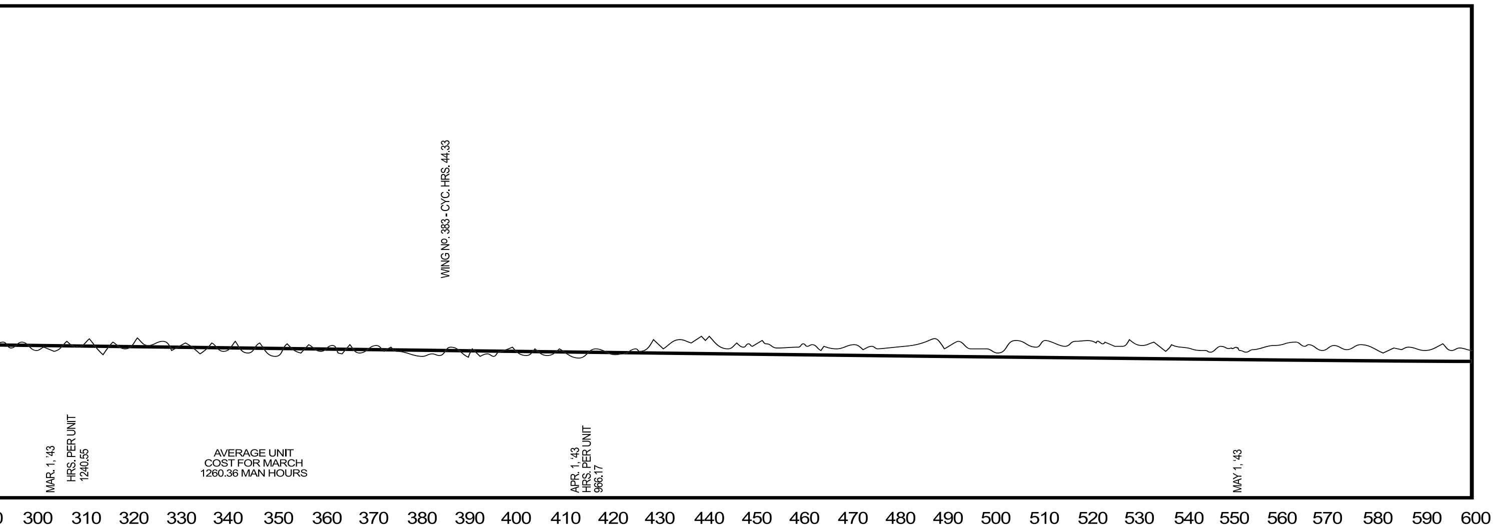

been written by a familiar composer.

A better-known power law example is Moore's Law for computing. Manufacturing has many familiar examples of power law relationships: The utilization of equipment in a job shop. The frequency of part usage in a part number database. The volume of sales of a series of models in a product line.

The power law expresses the "80-20 rule." Applied to quality defects, it is often asserted that 80 percent of the defects come from 20 percent of the known causes. Of course, 80-20 is not a precise ratio, but a way to indicate that defects follow a power law relationship. Eliminating a few major culprits eliminates a high percentage of defects. A Pareto chart roughly classifies observed data by a power law.

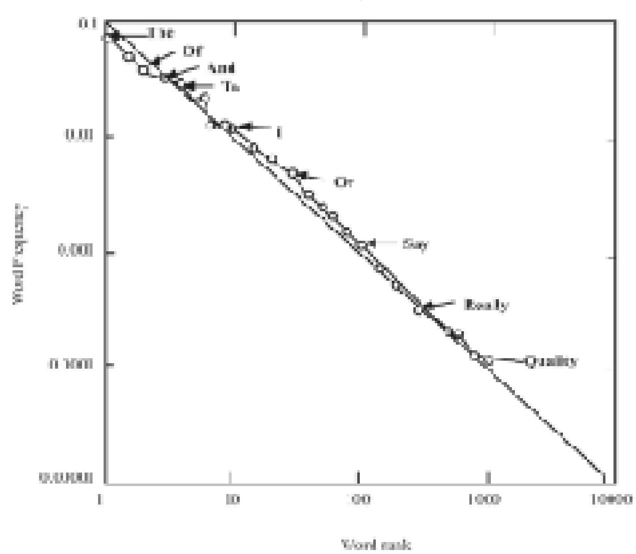

Figure 5. *Illustration of Power Law: "Zipf's Law" of frequency of word usage in the English Language. From* Would-Be Worlds, *John L. Casti, Santa Fe Institute, Santa Fe, NM, published by John Wiley, New York, 1997. Figure 3.14, p. 126.*

As with Zipf's Law, a power law relationship is thought to signify a system encoding complex information. A wide range of word frequencies make language interpretable by enabling complex, unique patterns. If every word were used with equal frequency, the message would be more like random noise than information. That is, a power law relationship suggests a system for meaningful information, but the relationship itself communicates no specific message, just as Zipf's Law describes no specific message in English. (All languages can be described by a power law.)

Sequences of codons in DNA follow a power law. Even "junk DNA" sequences follow a power law, which suggests that the junk is not devoid of useful information. However, no one yet knows what it is.

Some applications of the power law have made truly cosmic projections. For example, Richard Coren used power law relationships to project evolution over millennia, where evolution is ever-increasing "information content" in the universe, including human activity in our small corner of it.[4] Coren's projection is based on the observation that as they age, individual organisms (like people) have a longer time interval between significant learning events. However, a large population of interacting "organisms" consistently decreases the time between significant learning events. We change more slowly; our environment changes more rapidly.

Coren's lesson for process learning: It goes faster with genuine, coordinated, widespread human involvement by people who have the tools to do it.

A power law pattern indicates that a mass of detail is present, waiting to be investigated, but action is taken case-by-case. Just as with Zipf's Law for languages, little is learned staring at the graph. You must learn the specific language.

Perhaps Normality is an Illusion

For teaching the concept of emergence, a favorite power law example is avalanches. Huge, destructive avalanches command attention, but small avalanches occur with much greater frequency. Unless an observer is watching carefully, a lot of snow goes downhill unnoticed. Avalanche sizes and frequencies have been found to follow a power law.

Each avalanche occurs because a myriad of local conditions releases snow from its angle of repose. No two individual avalanches are exactly the same. One can collect approximate data and calculate the average frequency and average size of avalanches. However, reality is that no avalanche event is average.

A more obvious example is trying to calculate the average size branch on a tree. If one cut it up and measured weights of branches, the weight versus number would be a power law, but that's arbitrary. Where does one branch end and another begin? What of the little nubbins that have not yet clearly formed a branch?

Calculating the average branch diameter is even sillier. Each branch starts fat and disappears to nothing, or into other branches, and nowhere along the length is a branch's diameter uniform. Any average and standard deviation of tree branches is meaningless.

However, because most manufactured parts are expected to be nearly uniform – conform to a spec — seeing that each one, and the exact circumstances of its making, is unlike every other is much less obvious. Verifying such a thing would be a precision measurement project no one is likely to pay for, so we don't think about it. Historically, after design and ramp up, we only checked whether parts met spec. Quality much improved by thinking outside that box to compress process variance as much as possible – thinking "power law," not deviation from mean or from spec.

The argument has been made that years of using the normal distribution for statistical measurement has dulled our insight into the reality it supposedly represents, even in science. Many phenomena aren't statistically "normal" at all. We just think they are because our thinking follows the measurements we use.[5]

Quality thinking has begun to change "normal distribution" mindsets. We don't assume that it's enough to control an operation inside its normal limits. We don't wait to see real errors to take corrective action. In safety we take preventive action based on near misses. In quality, we do the same, codified in the phrase: "Quality is an attitude."

Out of the Box Organizational Performance

Considered organization-wide, out-of-the-box performance follows the analogy between the normal distribution and the power curve. Our legacy of business and managerial customs assumes that there is a "normal state" of operations. Once normal is achieved, "instinct" is for management to maintain that state. Any significant deviation from norm should take management permission.

Hiring practices, budgetary controls, supplier contracts, ad infinitum are regarded as control systems by caretaker managements. More aggressive managers set improvement goals, then control to meet them. At the company level, woe betides the management of a publicly traded firm that projects an earnings goal, but fails to meet it.

All these practices lend themselves to continue managing "in the box." Buried in our legacy of business and managerial customs is the assumption of "some normal state to be maintained" or "next normal state to be achieved." Just as with the old ramp-up learning curves, they continue to create mental boxes of our own making.

Perhaps that's the reason most companies pursuing a version of lean manufacturing stagnate at C-class as defined in Figure 6. Breaking out of that box is a total business, change – beyond making cultural modifications to improve flow and quality on the shop floor. By thinking of lean practices, including the human side of them, as techniques, a company can shuffle the organization. However, changing mindsets is a different game with a strange deck of cards.

The mind set change is subtle. Consider standard work according to the Toyota system, which sounds as boringly close to holding a norm as one can get.

Without adhering to method, you don't know "where you are" improving it. However, thought of as part of continuous improvement, or constant learning, standard work has feedback built into every work cycle: Finish early, why? Take extra time, why? What's happening outside the work center box that could be a cause? But beyond that, one cycle of standard work is never exactly like another, and like a power law language, actual performance reveals clues to making it "better." And better may be defined as handling more task variety within the work cycle, not just making sure that everything is acceptably the same.

Whether working line assembly or developing business unit strategy, the mindset to persist in that kind of detailed learning does not come naturally. A work group "really into the process" constantly sees ideas for improvement, noting the little avalanches, and not just the big ones. If improvement seldom occurs unless management or staff prompts it, the task of human development is not done.

How to organize for this? The typical organization chart is – well – a set of boxes that lend themselves to setting up budget controls. Reorganization usually shuffles boxes or shuffles people among them, as with a tinkertoy set. Knowing primary responsibilities is vital, but for superior per-

Summary of A-B-C (or C-B-A) Classifications

Pre-C	Experimentation. Try it; you'll like it. No serious organizational change. Results are fragmented. Little overall benefit.
C-Class	A comprehensive effort covering at least one complete operating organization. The focus is on technique adoption and development. Have results to demonstrate, and they may be substantial. Operational links to some customers and suppliers. Still directing change more than leading it. Sometimes distracted by competing initiatives or techniques. The cultural change is incomplete, so measurements and status systems may conflict with support of process compression. Not yet self-sustaining, so may be quickly blown away.
B-Class	The culture is maturing. The total unit organizes to support compression of value-adding processes. Changes are now more led than directed. Leaders emphasize vision, mission, goals, and values. Work more by "principles" than by emulation of techniques. Have more and closer linkages to customers and suppliers. Measurements and rewards support the system. Improvement is self-sustaining. Changing this culture would take time, so it's not easily dispelled. But if the organization becomes obsessed with process improvement, it may not be as innovative as it could be – lose sight of objectives of processes because of the dedication to their improvement.
A-Class	Both innovative and operationally efficient. Extend the "principles" or expand "excellence" across more goals and in new directions. Able to integrate new technology into the "excellence philosophy." May no longer be a single, conventional organization, but a flexible coalition of working partners. Experiments with new business models.

The same figure appeared in "The A-B-Cs of Excellence," *Target*, Vol. 17, No. 3, p. 12.

Figure 6.

formance, people in their boxes (or cubicles) see a long way outside them. Otherwise they strive to improve the performance of only their little box.

If status systems and rewards promote fragmentation, performance is the sum of the boxes. Organizations have tried all kinds of formulas to break this syndrome, but thinking of the problem structurally is still being in a box. Regarded only as the adoption of techniques, lean manufacturing remains limited to C-class performance. Improvement is a now and then thing, often initiated by a coordinator, to control what has been achieved. Managers rewarded for lean projects are likely to do the thinking and personally direct them rather spending their time stimulating others. Take away the champion, and lean will collapse.

Sustaining the improvement or going to B-Class, much less A-Class, is not the addition of more techniques or software. It's a change to a genuine learning culture, one that keeps learning curves going and going. Organizational changes and technique improvements follow naturally in the pursuit of the learning culture.

A Different Kind of Leadership

Leadership to change culture to a learning mindset uses a different language, leading by example and by asking questions, stimulating people to see and to think. Learning people sustain learning processes. As people develop, they begin to improve operations on a broader and broader scope. Teaching people to see and to think is not the exercise of control in some normal state. There is no such thing.

Without development of people as improvement agents, tools and organizational changes, while necessary, are of limited effectiveness. If one changes to a "product line" organization," for instance, people will just accommodate to new boxes unless they learn to see and work a long way outside of them.

From control systems to pay policies, the received tradition of the boxes militates against this. A few examples from Toyota may help. They do not motivate people primarily by money. They don't emphasize bonuses that reinforce status of the boxes, or of those within them. Despite ingenious devising, those are apt to result in unintended dysfunctional behavior, and at worst, set people on one another when they should be collaborating. Much the same can be said of company data where it is harbored and guarded by departments. People cannot learn in a closed information environment.

Toyota emphasizes a goal of vehicles that people will love. It stirs concern that a competitor might find a little edge somewhere. It stimulates tribal affiliation with Toyota and not for some box within it. And Toyota is far from breaking as many people out of their boxes as it would like.

Every aspect of an organization and its leaders' behavior affects how every person within it can work together to improve performance. It's impossible to "implement" a transition to achieve B-Class in the usual sense. Leaders have to learn to lead out-of-the-box, a radical departure from in-box legacies and instincts.

Footnotes

1. "The Manufacturing Progress Function," R.W. Conway and Andrew Schultz, Jr., *Journal of Industrial Engineering*, Vol. 10, No. 1, Jan.-Feb. 1959, pp. 39-54.

2. For discussion of the limits of models and the phenomena of mind set limitations, see "Out of the Box," Robert W. Hall, *Target*, Vol. 17, No. 2, 2001, pp.16-22.

3. "Comparative Review of Cost Reduction Programs in Manufacturing," Frank T. Lastra, *Journal of Industrial Engineering*, Vol. 16, No. 5, Sept.-Oct. 1965, pp. 340-345.

4. *The Evolutionary Trajectory*, Richard L. Coren, Gordon & Breach Science Publishers, 1998.

5. An example of such an argument is "Two Lessons from Fractals and Chaos," Larry S. Liebovitch and Barbara Scheurle, *Complexity*, Vol. 5, No. 4, March-April, 2000, pp. 34-43. (Published by John Wiley & Sons.) The tree limb example is from that article.

Robert W. Hall is editor-in-chief of Target *and a founding member of the Association for Manufacturing Excellence.*